Accelerating Our Understanding of Supernova Explosion Mechanism via Simulations and Visualizations with GenASiS

Reuben D. Budiardja
National Institute for Computational Sciences
The University of Tennessee
Knoxville, Tennessee
reubendb@utk.edu

Christian Y. Cardall
Physics Division, Oak Ridge National Laboratory
Oak Ridge, Tennessee
Department of Physics and Astronomy, University of Tennessee
Knoxville, Tennessee
cardallcy@ornl.gov

Eirik Endeve
Computer Science and Mathematics Division, Oak Ridge National Laboratory
Oak Ridge, Tennessee
Department of Physics and Astronomy, University of Tennessee
Knoxville, Tennessee
endevee@ornl.gov

arXiv:1509.08779v1 [astro-ph.IM] 29 Sep 2015

ABSTRACT

Core-collapse supernovae are among the most powerful explosions in the Universe, releasing about 10^{53} erg of energy on timescales of a few tens of seconds. These explosion events are also responsible for the production and dissemination of most of the heavy elements, making life as we know it possible. Yet exactly how they work is still unresolved. One reason for this is the sheer complexity and cost of a self-consistent, multi-physics, and multi-dimensional core-collapse supernova simulation, which is impractical, and often impossible, even on the largest supercomputers we have available today. To advance our understanding we instead must often use simplified models, teasing out the most important ingredients for successful explosions, while helping us to interpret results from higher fidelity multi-physics models. In this paper we investigate the role of instabilities in the core-collapse supernova environment. We present here simulation and visualization results produced by our code GenASiS.

1. INTRODUCTION

Core-collapse supernovae are the violent deaths of massive stars and are among the most powerful explosions in universe, releasing about 10^{53} erg of energy on timescales of a few tens of seconds. This rivals the instantaneous power of all the rest of the luminous visible universe combined. They mark the birth of the most exotic states of matter known: neutron stars and black holes, while at the same time producing and disseminating most of the elements heavier than helium, making life as we know it possible. These events occur about twice per century in a typical galaxy like our own and have been in the forefront of research in the field of astronomy and astrophysics for almost half a century. Yet how exactly they work is still shrouded in mystery.

Stars burn hydrogen into helium for most of their existence. For stars more massive than ~ 10 solar masses ($M_\odot$), temperatures and densities are sufficiently high for burning to continue through carbon and to iron group elements. The star ends up in an onion-like configuration, with an iron core surrounded by layers of silicon, oxygen, carbon, helium, and hydrogen. Since the iron group elements are the most tightly bound, burning in the core ceases. At this point, the pressure in the core is dominated by electron degeneracy pressure (a consequence of the Pauli exclusion principle), which supports it against the inward pull of gravity. This balance between the gravitational pull and the electron degeneracy pressure in the core is only marginally stable.

Two processes occur in the core that result in the reduction of the degeneracy pressure support: electron capture on the free protons and nuclei, and nuclear dissociation under extreme densities and temperatures. The pressure support in the core is reduced enough that the core eventually becomes unstable and collapses. As the core collapses, the inner and outer regions behave differently. The inner core undergoes homologous collapse—velocity increases linearly with radius—as expected of a fluid with relativistic, degenerate electron pressure. With increasing radius, the density decreases, and thus also the local sound speed. Thus, there is a radius where the speed of the infalling matter is the same as the local sound speed, demarcating the inner and outer core. Beyond this radius—the outer core—matter collapses supersonically.

The inner core collapses until it exceeds nuclear matter density ($\sim 1-3\times10^{14}$ g/cm^3). At this extreme density, the pressure of the inner core increases dramatically as a result of the repulsive component of the short-range nuclear force. The inner core becomes incompressible and bounces, and a shock wave forms at the boundary of the inner and outer core and begins to move out. Ultimately this shock wave will be responsible for the disruption of the star, producing

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

XSEDE '15, July 26 - 30, 2015, St. Louis, MO, USA

© 2015 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-3720-5/15/07…$15.00
DOI: http://dx.doi.org/10.1145/2792745.2792746

the observable explosion.

It was once thought that as the shock wave propagated outward, the velocity of the bounce would grow as it moved into the outer layers of the core; the bounce would therefore be the origin of the supernova's energy [1]. From all the more realistic models completed to date, we now know that this is not the case, and therein lies the core-collapse supernova problem.

As the shock propagates out, it has to move through infalling material in the outer core, during which nuclear dissociation happens. This costs the shock energy. Additional energy losses occur when electron capture on the free protons liberated by nuclear dissociation eventually results in an electron neutrino burst. As a result of these energy losses, the shock stalls.

If this were the end of the story, no supernova would ever explode nor be observed. The shock has to be reenergized so that it may continue to propagate outward and eventually produce the explosion. The details of how the stalled shock is revived is the central question in core-collapse supernova theory.

Out of the 10^{53} erg released during the explosion, the visible explosion energy is only 1%. The rest is released as neutrinos. Because neutrinos dominate the energetics of a supernova event, it is natural to consider neutrino heating as a mechanism for the revival of the stalled shock. This delayed neutrino-heating has been proposed as one of the mechanisms that leads to explosion (for example, see [2] for reviews).

Core-collapse supernovae are asymmetric events. Observational evidence that has accumulated to support this includes spectropolarimetry, large average pulsar velocities, and the morphology of highly resolved images of supernova such as SN 1987A (see [3] and references therein). On the theoretical side, simulations have shown that a variety of fluid instabilities are present. These instabilities develop convective overturn and help transport hot gas from neutrino-heating region directly to the shock, thereby enhancing the neutrino energy deposition to the stalled shock (e.g [4] and references). These multidimensional effects therefore may be important for the neutrino-heating mechanism to revive the stalled shock. Recent simulations have also revealed the existence of the standing accretion shock instability (SASI), which given enough time, may also grow and impact the dynamics [5]. All these multidimensional effects may play essential roles in possible mechanisms of core-collapse supernovae.

Because of the complexity of the candidate supernova mechanisms, a purely analytic investigation is not possible. Instead, supernova modeling requires sophisticated numerical simulations. All the input physics required to model core-collapse supernovae present daunting challenges that are both algorithmic and computational in nature, and will tax state-of-the-art supercomputers for years to come. This is especially true for three-dimensional models and simulations, where multi-physics simulations with the required level of physical fidelity will not fit even the largest supercomputers currently available. Instead, in three spatial dimensions we must often resolve to simplified models to tease out the important ingredients for successful supernova explosions. There is indeed a value to this approach as it also helps us to interpret results from the much more computationally expensive multi-physics models.

The investigation reported in this paper follows the same approach. In particular, we are interested to learn more about the role of hydrodynamics instabilities in core-collapse supernova environments. We build on previous results by [6] and extend their work to three spatial dimensions. This paper is organized as follows. We first describe our numerical code in section 2. In section 3 we describe our simulations and results, followed by summary and conclusion in section 4. All of our simulations were done using the Darter supercomputer [7] at the National Institute for Computational Sciences (NICS).

2. THE TOOL: GENASIS

GENASIS (*Gen*eral *A*strophysical *Si*mulation *S*ystem) is a new code we are developing to facilitate simulations of astrophysical problems on high-performance computing resources. The word 'General' in its name denotes its extensibility to implement and use multiple physics, solvers, and numerical algorithms by virtue of abstracted names and/or interfaces. Rather than a single program, it is a 'Simulation System' comprising a collection of modules, structured as classes, which suitable driver programs can use to initialize and solve particular problems of interest. 'Astrophysical' denotes the kinds of applications the code is ultimately aimed at and the types of solvers it will eventually provide. In this section we briefly discuss some features of GENASIS relevant to the simulations at hand.

2.1 Numerical Hydrodynamics

The equations of ideal Newtonian hydrodynamics in conservative Eulerian formulation take the form

$$\frac{\partial D}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^i\right) = 0, \quad (1)$$

$$\frac{\partial S^j}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^j v^i + p\delta^{ij}\right) = \mathcal{S}_M^j, \quad (2)$$

$$\frac{\partial E}{\partial t} + \frac{\partial}{\partial x^i}\left(\left[e + p + \frac{1}{2}\rho v^j v^j\right] v^i\right) = \mathcal{S}_E. \quad (3)$$

These are the equations of mass, momentum, and energy conservation, respectively. In these equations, ρ, v^i, e, and p represent rest mass density, fluid velocity, internal energy density, and fluid pressure, while $\mathcal{S}_M^j$ and $\mathcal{S}_E$ represent momentum and energy source terms. Summation over repeated indices is implied. The time-evolved *conserved* quantities D, E, and S^j are related to the *comoving* variables ρ, e, v^i through the relations

$$D = \rho, \quad (4)$$

$$S^j = \rho v^j, \quad (5)$$

$$E = e + \frac{1}{2}\rho v^j v^j. \quad (6)$$

The system of equations is closed by an equation of state relating the pressure to the internal energy. For an ideal gas, a polytropic equation of state of the form

$$p = e(\Gamma - 1) \quad (7)$$

is used where Γ is the adiabatic index.

In GENASIS we have implemented a solver for Equations (1) - (3) based on the finite volume method. This is done by casting the equations as a system hyperbolic conservation

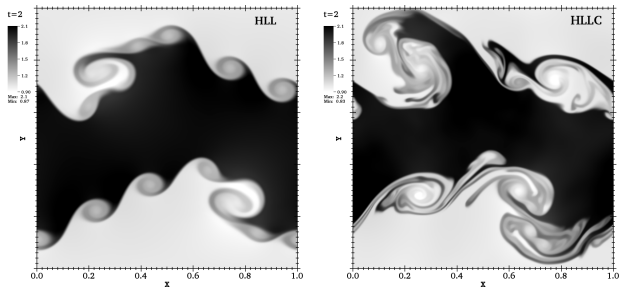

Figure 1: Density plots of Kelvin-Helmholtz instability test as computed using HLL (left) and HLLC (right) solver on 512^2 grid. The HLLC solver produces a more accurate and complex structure of the instability.

laws of the form

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f} = 0, \tag{8}$$

where the vector of the conserved variables $\mathbf{u}$ and their corresponding fluxes $\mathbf{f}^i$ are defined as

$$\mathbf{u} = \left[D, S^j, E\right], \tag{9}$$

$$\mathbf{f}^i = \left[\rho v^i, \rho v^j v^i + p\delta^{ij}, \left(e + p + \frac{1}{2}\rho v^j v^j\right) v^i\right]. \tag{10}$$

The fluxes through cell faces are computed using the so-called HLL-type approximate Riemann solvers [8][9]. We have implemented two variants of these solvers. The first one, which we simply label 'HLL'—for Harten, Lax, van Leer who first devised the method [8]—is given by the expression

$$\mathbf{f}^{\mathrm{HLL}} = \frac{\alpha^+ \mathbf{f}^{\mathrm{L}} + \alpha^- \mathbf{f}^{\mathrm{R}} - \alpha^+ \alpha^- \left(\mathbf{u}^{\mathrm{R}} - \mathbf{u}^{\mathrm{L}}\right)}{\alpha^+ + \alpha^-}. \tag{11}$$

The superscripts L and R denote the values on the left and right sides of the cell face, respectively, as given by equations (9) - (10). The coefficients $\alpha^{\pm}$ are constructed from the maximum and minimum characteristic speeds $\lambda^{\pm}$ as

$$\alpha^{\pm} = \max\left\{0, \pm\lambda^{\pm}\left(\mathbf{v}^{\mathrm{L}}\right), \pm\lambda^{\pm}\left(\mathbf{v}^{\mathrm{R}}\right)\right\}, \tag{12}$$

where

$$\lambda^{\pm} = v \pm \sqrt{\Gamma p/\rho}. \tag{13}$$

The second HLL Riemann solver variant we have implemented, known as the HLLC solver [9], takes into account both the acoustic and entropy waves which results in higher accuracy in preserving contact discontinuities. Our implementation of the HLLC solver includes a fallback to the more dissipative HLL solver in grid-aligned shocks to avoid "odd-even decoupling" [10]. We have verified our implementation of these solvers using various numerical test problems with known solutions. These are detailed in [11]. Figure 1 shows the effect of the HLLC solver relative to the HLL solver on the Kevin-Helmholtz instability test problem.

In GenASiS the conserved and comoving variables are defined as cell-centered values and must be reconstructed on the cell faces for the flux computation. We use slope-limited linear interpolation (which for cells of equal size yields a spatially second-order accurate scheme for smooth flows) to preserve non-oscillatory behavior near shocks and discontinuities. Specifically we have implemented a one-parameter ($\theta \in [1, 2]$) family of generalized minmod limiters [12], which for an arbitrary variable χ is given by

$$\left.\frac{\partial \chi}{\partial x}\right|_{ijk} = \mathrm{minmod}\left[\begin{array}{l} \theta\left(\frac{\chi_{ijk} - \chi_{i-1jk}}{x_i - x_{i-1}}\right), \\ \left(\frac{\chi_{i+1jk} - \chi_{i-1jk}}{x_{i+1} - x_{i-1}}\right), \\ \theta\left(\frac{\chi_{i+1jk} - \chi_{ijk}}{x_{i+1} - x_i}\right) \end{array}\right] \tag{14}$$

where the multi-variable minmod function is defined as

$$\mathrm{minmod}\left(\chi_1, \chi_2, \dots\right) = \begin{cases} \min_j\left(\chi_j\right), & \text{if } \chi_j > 0 \ \forall j, \\ \max_j\left(\chi_j\right), & \text{if } \chi_j < 0 \ \forall j, \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

In the simulations presented here we have used $\theta = 1.4$.

For time-stepping, we use a second-order TVD Runge-Kutta scheme [13]:

$$u^{(1)} = u^n + \Delta t L(u^n) \tag{16}$$

$$u^{(n+1)} = \frac{1}{2}u^n + \frac{1}{2}u^{(1)} + \frac{1}{2}\Delta t L(u^{(1)}), \tag{17}$$

where the operator $L(u^n)$ denotes the spatial differencing. The time interval Δt must obey the Courant-ĂŞFriedrichsĂŞ-Lewy (CFL) condition given by

$$v_c \cdot \Delta t < C\Delta x, \tag{18}$$

where v_c is the maximum characteristic speed on the grid, Δx is the cell spacing, and $C < 1$ is the "Courant parameter."

2.2 Multilevel Mesh Refinement

Adaptive mesh refinement (AMR) is a technique to dynamically furnish higher resolution only when and where the problem requires it, thereby enabling focused deployment of limited computational resources. This technique becomes increasingly important in multi-physics codes since advances in scientific modeling, algorithms, and physical fidelity have so far always exceeded resources that can be provided even by the largest supercomputers.

There are two common approaches to AMR. Block-structured AMR [14], as the name implies, uses blocks of cells as the basic unit of the mesh. The coarsest grid, consisting of several blocks, covers the entire computational domain. Each block may be refined, creating finer nested grids.

In developing GenASiS however we have instead chosen to use cell-by-cell AMR [15]. In this approach, individual cells may be refined (and coarsened) as needed; this fine-grained control leads to a smaller number of total cells and therefore potentially larger savings of computational resources. Our motivation for this approach is based on the very high computational cost for each spatial cell in the types of models we eventually intend to run. For example, in its full glory the solver for the neutrino radiation transport equations will have to compute the coupled interactions in momentum space for every spatial cell. In this case, the flexibility of cell-by-cell AMR may prove to be advantageous overall.

In our AMR implementation, we take a level-by-level approach in arranging storage and solvers. Each level of refinement is a mesh in its own right. The meshes at different

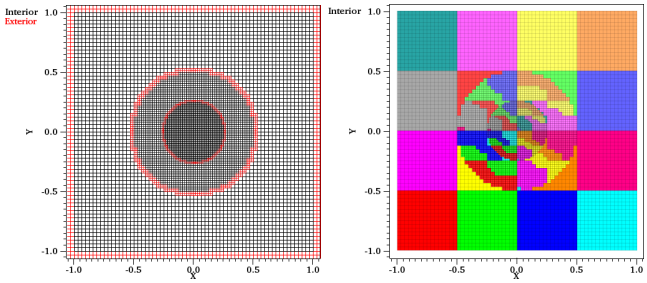

Figure 2: Illustrations of a multilevel mesh with refinement for a total of three levels (left) and their domain decomposition (right) where different processes are indicated by different colors. The red cells (left) show the exteriors of the level meshes, which is the computational boundary at the coarsest level and the coarse/fine boundary on the finer levels. On the finer levels, the exterior values are prolongated from the coarser levels.

levels of refinement are therefore domain-decomposed independently for parallel processing (see Figure 2 for illustration). This approach simplifies the formulation of solvers since they only need to be written for a single level, while interactions between levels are handled separately. In our level-by-level approach, solutions from finer levels are used to improve solutions on coarser levels via restriction operations, and solutions from coarser levels are used as boundary conditions on finer levels via prolongation operations. Further details of our level-by-level formulation are given in [11].

2.3 Parallelization

GenASiS uses a simple 'brick' decomposition to domain-decompose the computational domain into several subdomains. The brick decomposition simply means that, in three dimensions, the computational domain is divided in each dimension by $n_b = \sqrt[3]{n_p}$, i.e. the cube root of the number of processors n_p. Each subdomain is then assigned to an MPI process.

For solvers requiring nearest-neighbor cell values crossing process boundaries, each subdomain keeps 'ghost' cells in addition to the 'proper' cells (the normal working computational cells assigned to the process). The values in ghost cells are then populated via point-to-point message passing with neighboring processes that own the corresponding proper cells. The brick decomposition maximizes the ratio of proper-to-ghost cells, therefore minimizing the nearest-neighbor communications.

To facilitate these exchanges of ghost cell data, we created two sets of "connectivity" for the proper cells. The "exchange" connectivity contains proper cells whose values need to be communicated to neighboring processes to populate their ghost cells, while the "non-exchange" connectivity contains the rest. Required computations are performed first on the exchange cells, after which a non-blocking, point-to-point MPI communication is initiated. While communication continues in the background, computations are performed on the non-exchange cells. This technique of overlapping work and communication proves to be essential to achieve the high-scalability of our code, as shown in Figure 3.

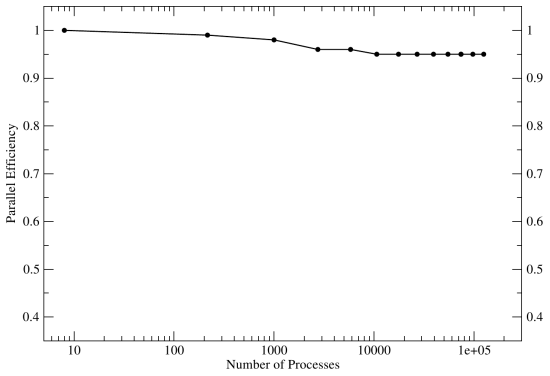

Figure 3: MPI weak scaling of the hydrodynamics solver in GenASiS with a single-level mesh, where the number of computational cells per MPI process is kept fixed to 48^3 as the number of processes are increased. High efficiency of around 0.95 is maintained up to $100,000$ MPI processes.

3. THE INVESTIGATION: THE ROLE OF CONVECTION AND SASI

3.1 Initial Conditions

We adopt a simplified model of the post-bounce core-collapse supernova environment in three dimensions. We use an ideal gas equation of state (c.f. equation 7) with $\Gamma = 4/3$. Our model accounts for (momentum and energy) sources due to central gravity, heating and cooling due to neutrino radiation (Q_ν), and nuclear dissociation/recombination (Q_D) as

$$\mathcal{S}^j_M = -\rho \frac{\partial \Phi}{\partial x^j} \tag{19}$$

$$\mathcal{S}_E = -\rho v^i \frac{\partial \Phi}{\partial x^i} + Q_\nu + Q_D, \tag{20}$$

where $\mathcal{S}^j_M$ and $\mathcal{S}_E$ enter into the hydrodynamics equations as given by Equations (2)-(3). We use the point mass gravitational potential $\Phi = -GM/r$ where G is the gravitational constant, M is the mass of the central compact object (i.e. the proto-neutron star), and r is the radial distance from the center of the star. Throughout our simulations we set $M = 1.3\ M_\odot$, where $M_\odot$ is the Solar mass.

We follow [6] in parameterizing the heating and cooling effects due to neutrino radiation, using values for the parameters that roughly mimic those in self-consistent supernova simulations. The heating and cooling function is given by

$$Q_\nu = \rho \left(\frac{B}{r^2} - A p^{3/2} \right) e^{-(s/s_0)^2} H(\mathcal{M}_0 - \mathcal{M}), \tag{21}$$

where B is a parameter proportional to the fixed neutrino luminosity (i.e. heating) as derived in [16]. The values for the cooling parameter A are determined empirically as those which yield zero fluid velocity at the radius of the proto-neutron star $R_{\rm PNS}$ for a shock at radius $R_{\rm sh} = 100$ km with no heating (i.e $B = 0$). We set $R_{\rm PNS} = 40$ km in all of our simulations as the fixed inner boundary. $s \propto \log(p/\rho^\Gamma)$ is the entropy with the referential value s_0 for the minimum initial post-shock entropy. The Heaviside step function $H(\mathcal{M}_0 - \mathcal{M})$ excludes heating and cooling outside of the shock based on the Mach number $\mathcal{M}$ with $\mathcal{M}_0 = 2$.

We include a nuclear dissociation energy per unit mass ϵ in the shock jump conditions which in turn influences whether the instability is *SASI-dominated* or *convection-dominated*. Larger values of ϵ corresponds to larger shock compression ratios and slower post-shock speeds. For these simulations, in addition to equations (1)-(3) we also solve for the mass density of the dissociated baryons ρ_D:

$$\frac{\partial \rho_D}{\partial t} + \nabla \cdot (\rho_D \mathbf{v}) = R_D, \tag{22}$$

where the source term R_D is given by

$$R_D = \frac{\rho - \rho_D}{\Delta t} \quad \text{(dissociation)}, \tag{23}$$

$$= -\frac{\rho_D}{\Delta t} \quad \text{(recombination)}. \tag{24}$$

For numerical stability we spread the change over several time steps by taking $R_D \rightarrow f R_D$, where $f = \min(0.1, f_e)$, in which f_e is the value that yields a 1% change in internal energy density due to dissociation. The associated energy source term due to dissociation is then

$$Q_D = -\epsilon R_D. \tag{25}$$

3.2 Simulation Setup

We produced two series of models: a SASI-dominated series with $\epsilon = 0$ (no dissociation), and a convection-dominated series with $\epsilon = 0.3$. For each series we increased the heating value B until explosions were attained. For each value of B the initial shock radius $R_{\rm sh}$ was determined as that which yielded zero fluid velocity at $R_{\rm PNS}$ using the previously obtained (with zero heating) value of A. We used Mathematica[1] to obtain initial profiles, solving spherically symmetric steady-state versions of equations (1)-(3) with the Rankine-Hugoniot shock jump conditions joining the pre- and post-shock regions.

The initial profile was then mapped to our three-dimensional multi-level mesh with refinement, covering the computational domain $[-640, 640]$ km. A total of four mesh levels was used. The coarsest level mesh had a resolution of 128^3 covering the entire computational domain, yielding a spatial resolution of 10 km. Each subsequent level increased the resolution by a factor of two in each dimension. The finest level of mesh therefore had a resolution of 1.25 km. 64 MPI processes were used to run each model in parallel. The computational cost for a single model was about $12,000$ SU on Darter.

The initial condition was perturbed by introducing random pressure perturbations of amplitude 0.1% relative to the local value. The perturbations seeded the instabilities. For each value of heating for each series, we ran ten simulations of the same model (differing only by their initial random perturbations). We let each simulation run to either 1 second of physical time or until explosion was attained.

Most of the simulations presented here were run in the last few weeks of the year 2014, taking advantage of Darter's lower utilization rate due to the holiday. To get better job throughput for our simulations, and for easier job management, we bundled the ten simulations in each model into a single job request on Darter. Furthermore, when computational resources were available, we spread out the rank placement of our jobs, using only four to six MPI processes per socket. This provided higher memory bandwidth and effectively boosted the computational speed by 30 to 40 percent, since the hydrodynamics algorithm is inherently memory-bandwidth bound.

[1] http://www.wolfram.com/mathematica/ version 10.0

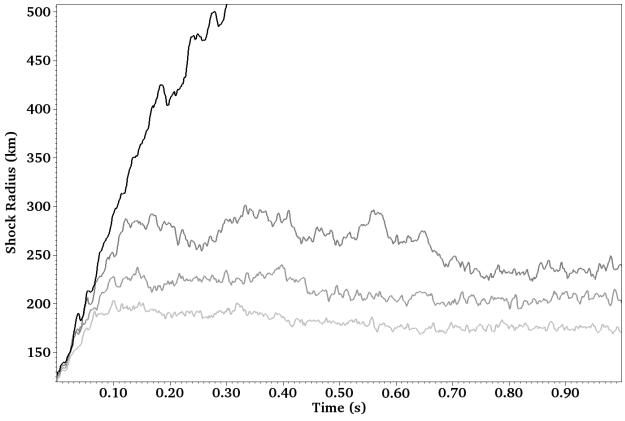

Figure 4: Shock trajectories of representative convection-dominated models with heating values $B = [0.65,\ 0.70,\ 0.75,\ 0.80]$, where darker color indicates higher B. Only one of the ten simulations for each value B is plotted.

3.3 Results

For the series of convection-dominated models ($\epsilon = 0.3$), as we increase the heating values B the shock expands and settles at larger radius. Figure 4 shows the shock radii of some representative convection-dominated models using various heating values B. For these models, there seems to be a critical heating value B for which explosions are attained. In our simulations this value is between $B = 0.75$ and $B = 0.80$. All the convection-dominated models at and below $B = 0.75$ did not attain explosion, while all the models at and above $B = 0.80$ attained explosions. Figure 5 shows the shock radii of the convection-dominated models for $B = 0.80$. Although the shock radii do not follow exactly the same trajectories (due to different perturbations of the initial conditions), all of the simulations for this model have roughly similar explosion time scales.

The series of SASI-dominated models tells a completely different story. For these models, there does not seem to be a critical heating value B for which explosions are guaranteed. Rather, as we increase the heating values B, *only the probability of explosion increases*. Figure 6 plots the shock radii of ten SASI-dominated simulations for the model with heating value $B = 1.025$. Exploding simulations are plotted with dash-lines, while non-exploding ones are plotted with solid lines of different colors to indicate the different simulations. In the upper panel of Figure 6, a *zoomed-in* version of the plot is shown for the first 200 milliseconds of the simulations. The only differences in these simulations are the small random perturbations in their initial conditions. As seen on the upper panel of Figure 6, for at least the first 50 milliseconds, all simulations follow a very similar path in their shock radius plots. By 100 milliseconds, the simulations have already begun diverging in their evolution. In the lower panel of Figure 6, we see that eventually four simulations attained explosions for this heating value while the rest did

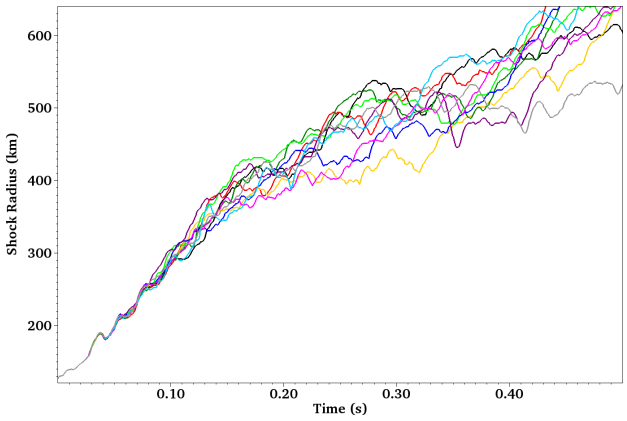

Figure 5: Shock trajectories of ten simulations (indicated by different colors) at heating value $B = 0.80$ for the convection-dominated model.

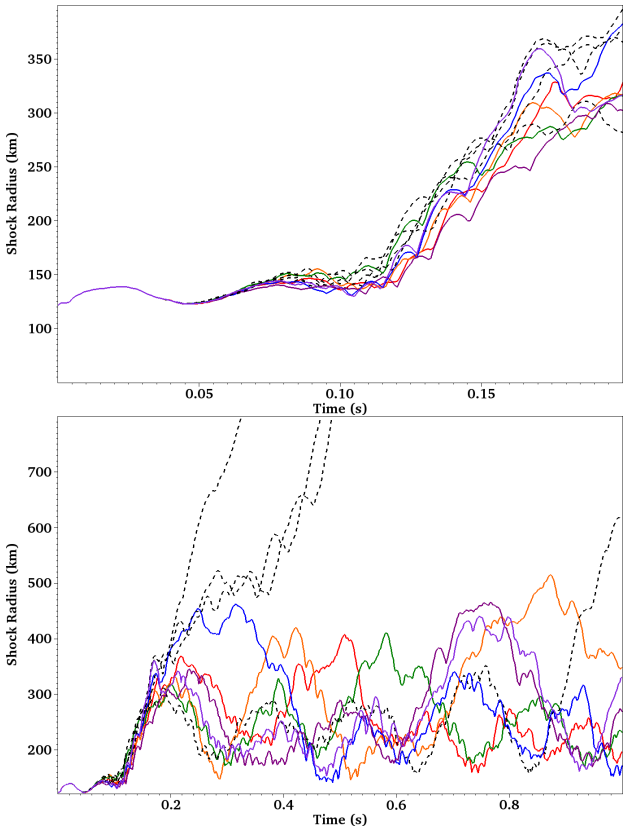

Figure 6: Shock trajectories of ten simulations for model at heating value $B = 1.025$ for SASI-dominated series. The upper panel shows a zoom-in of the trajectories at early time.

not. Furthermore, for the exploding simulations, explosions were reached at vastly different time scales: one exploding simulation reached critical shock radius of 600 kilometers at around 250 milliseconds, two exploding simulations reached it at around 500 milliseconds while maintaining much of the oscillatory behavior characteristic of SASI, and one simulation reached it at a much later time of nearly one second of evolution after a long period of vigorous SASI activity.

3.4 Visualizing the Data

Scientific visualization is an integral part of many scientific endeavors, helping us glean valuable first insights from data generated by the simulations. The higher dimensionality nature of the supernova problem in particular—three dimensions in position space plus another three dimensions in momentum space in radiation transport—and the size of the data from each simulation often push the innovation and capability of scientific visualization. For the simulations discussed here however, although each simulation produces only a modest amount of data (order of 2 Terabytes per simulation), the number of simulations done and the number of images and movies generated require us to develop and automate some kind of visualization workflow, which we discuss here.

GenASiS has interfaces to the Silo library[2] to be used for both checkpoint-restart files and output files. Silo was chosen for its portability and because it is readily readable by the visualization package VisIt,[3] which we have been using extensively.

We use entropy to visualize the fluid, with higher entropy indicating hotter fluid. To help visualize the three-dimensionality of the problem, two-dimensional slices parallel to the $xy-$, $xz-$, and $yz-$planes through the center of the computational domain are projected to the corresponding rear walls of the box bounding the computational domain. We accomplish this by using the `Slice` operator in VisIt to get the two-dimensional slices and the `Transform` operator to put the slices on the walls of the bounding box.

To show the shock surface, we use a `Contour` plot of the quantity MachNumber $= 2.0$. We set the opacity of this plot to 25 percent to allow us to see through the shock surface. We also show the isosurface of entropy value 0.15 using the `Contour` plot. These two plots display the three-dimensionality of the shock structure from a fixed viewing angle.

For every output of a single time slice from the simulations, a visualization frame with the above plots is created. Multiple frames are then assembled to make a visualization animation (e.g. movie) from a single simulation. To show progress through the animation, and to give a better sense of the volume of the shocked material, we also plot shock volume averages using 1D `Curve` plots, with solid lines to indicate the shock trajectory thus far in the simulation and dotted lines to indicate the future trajectory. The two frames (entropy/fluid plots and shock volume Curve plot) for each time slice are then merged using the image processing tool ImageMagick.[4]

We used supercomputing resources Darter and Beacon at NICS to create the visualization movies from the simula-

[2]https://wci.llnl.gov/simulation/computer-codes/silo
[3]https://wci.llnl.gov/simulation/computer-codes/visit/
[4]http://www.imagemagick.org/

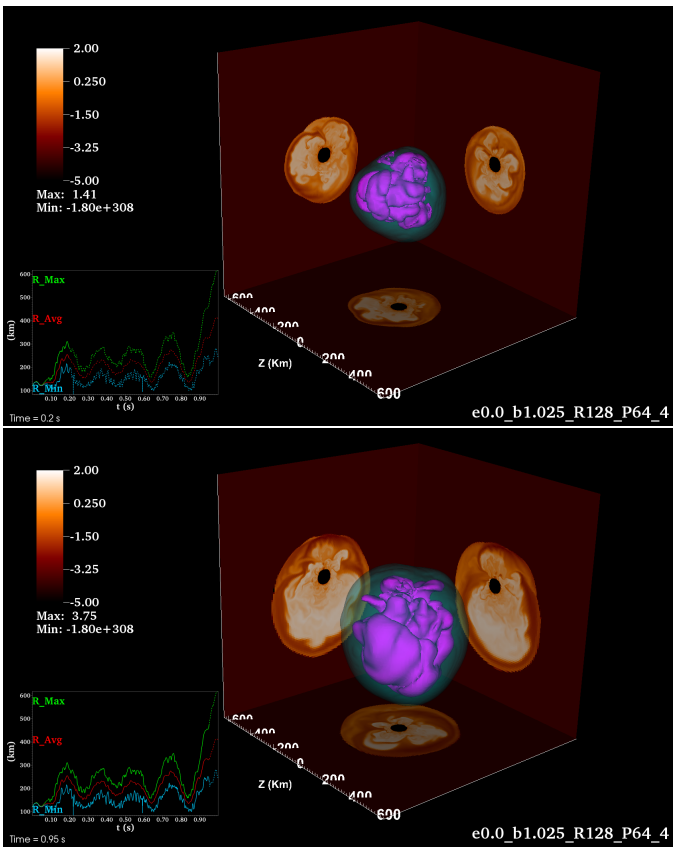

Figure 7: SASI-dominated simulation for $B = 1.025$ at early time $t = 0.2$ s and late time $t = 0.95$ s preceding explosion.

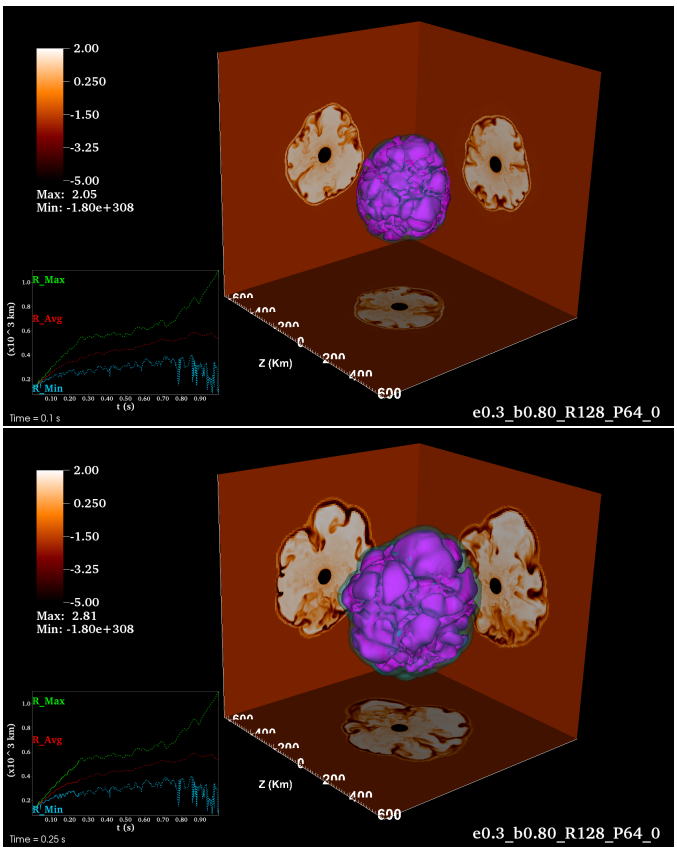

Figure 8: Convection-dominated simulation for $B = 0.80$ at early time $t = 0.1$ s and late time $t = 0.25$ s preceding explosion.

tions. Visualization scripts written in Python were used to drive VisIt and generate the plots described above. The VisIt scripts were then run in batch mode on Darter to do parallel rendering and generate PNG frames. Beacon was used to do image processing of the frames and to assemble them into movies using Mplayer software[5]. Beacon's large memory capacity and on-node solid-state drives were useful to process the images efficiently. NICS' system-wide Medusa filesystem was very useful in allowing us to use multiple resources for simulations, post-processing, and visualizations without the need to move large amounts of data.

Figures 7 and 8 show the visualization described here for one of the simulations at heating value $B = 1.025$ for the SASI-dominated model and $B = 0.80$ for the convection dominated model. Some visualization movies from the simulations will be presented during the conference.

4. CONCLUSION

In this paper we have presented an abridged description of our simulation results from over a hundred simulations of both convection- and SASI-dominated dynamics in the core-collapse supernova environment with GenASiS. A particularly important, previously undiscovered, result is the stochastic behavior of the SASI-dominated models. Because three-dimensional supernova simulations are typically very computationally expensive, it is unheard of to run tens, let alone hundreds of three-dimensional simulations. We were able to do so due to the simplified physics in our setup and the use of multilevel mesh refinement. Yet these models help us to tease out important ingredients of successful supernova explosions. They may also help us interpret results from more sophisticated physical models. Much more detailed analysis from our simulation results is currently still ongoing, and will be the subject of future studies.

Acknowledgment

This material is based upon work performed using computational resources Darter [7] supported by the University of Tennessee and Oak Ridge National Laboratory's Joint Institute for Computational Sciences (http://www.jics.utk.edu). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the University of Tennessee, Oak Ridge National Laboratory, or the Joint Institute for Computational Sciences.

5. REFERENCES

[1] Stirling Colgate and Montgomery Johnson. Hydrodynamic Origin of Cosmic Rays. *Physical Review Letters*, 5(6):235–238, September 1960.

[2] Hans-Thomas Janka. Explosion Mechanisms of Core-Collapse Supernovae. *Annual Review of Nuclear and Particle Science*, 62:407–451, 2012.

[3] R McCray. Supernova 1987A Revisited. *Annual Review of Astronomy and Astrophysics*, 31(1):175–216, September 1993.

[4] Anthony Mezzacappa. Ascertaining The Core Collapse Supernova Mechanism: The State of the Art and the Road Ahead. *Annual Review of Nuclear and Particle Science*, 55(1):467–515, 2005.

[5]https://www.mplayerhq.hu/

[5] John M. Blondin and Anthony Mezzacappa. The Spherical Accretion Shock Instability in the Linear Regime. *The Astrophysical Journal*, 642(1):401–409, 2006.

[6] Rodrigo Fernández, Bernhard Müller, Thierry Foglizzo, and Hans Thomas Janka. Characterizing SASI- and convection-dominated core-collapse supernova explosions in two dimensions. *Monthly Notices of the Royal Astronomical Society*, 440:2763–2780, 2014.

[7] Mark R. Fahey, Reuben Budiardja, Lonnie Crosby, and Stephen McNally. Deploying darter - a cray xc30 system. In JulianMartin Kunkel, Thomas Ludwig, and HansWerner Meuer, editors, *Supercomputing*, volume 8488 of *Lecture Notes in Computer Science*, pages 430–439. Springer International Publishing, 2014.

[8] Amiram Harten, Peter D. Lax, and Bram Van Leer. On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review*, 25(1):35, 1983.

[9] E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, July 1994.

[10] James J. Quirk. A contribution to the great Riemann solver debate. *International Journal for Numerical Methods in Fluids*, 18(6):555–574, March 1994.

[11] Christian Y Cardall, Reuben D Budiardja, Eirik Endeve, and Anthony Mezzacappa. GenASiS: General Astrophysical Simulation System. I. Refinable Mesh And Nonrelativistic Hydrodynamics. *The Astrophysical Journal Supplement Series*, 210:17, 2014.

[12] A Kurganov and E Tadmor. New High-Resolution Central Schemes for Nonlinear Conservation Laws and ConvectionâĂŞDiffusion Equations. *Journal of Computational Physics*, 160(1):241–282, 2000.

[13] C Shu and S Osher. Efficient implementation of essentially non-oscillatory shock-capturing schemes. *Journal of Computational Physics*, 77(2):439–471, August 1988.

[14] M. Berger and P. Collela. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64–84, May 1989.

[15] A Khokhlov. Fully Threaded Tree Algorithms for Adaptive Refinement Fluid Dynamics Simulations. *Journal of Computational Physics*, 143(2):519–543, 1998.

[16] H. Th. Janka. Conditions for Shock Revival by Neutrino Heating in Core-Collapse Supernovae. page 36, 2000.